arXiv:1611.06284v2 [cs.CV] 22 Nov 2016

Understanding Anatomy Classification Through Visualization

Devinder Kumar *
Philips Research
Eindhoven, Netherlands
devinder.kumar@philips.com

Vlado Menkovski †
Technische Universiteit Eindhoven
Eindhoven, Netherlands
v.menkovski@tue.nl

Abstract

One of the main challenges for broad adoption of deep convolutional neural network (DCNN) models is the lack of understanding of their decision process. In many applications a simpler less capable model that can be easily understood is favorable to a black-box model that has superior performance. In this paper, we present an approach for designing DCNN models based on visualization of the internal activations of the model. We visualize the model's response using fractional stride convolution technique and compare the results with known imaging landmarks from the medical literature. We show that sufficiently deep and capable models can be successfully trained to use the same medical landmarks a human expert would use. The presented approach allows for communicating the model decision process well, but also offers insight towards detecting biases.

1 Introduction

Understanding the decision process of a deep neural network model for classification can be challenging due to the very large number of parameters and model's tendency to represent the information internally in a distributed manner. Distributed representations have significant advantages for the capability of the model to generalize well [4], however the trade-off is the difficulty in communicating the model's reasoning. In other words, it is difficult to represent what information was used by the model and how to come to its particular output. In certain applications such as healthcare, understanding the decision process of a model can be a vital requirement.

One direction towards understanding how Convolutional Neural Netowrks (CNN) processes the information internally is through visualization. The work of Zeiler et. al [5], Mahendran et. al [1], and Zintgraf et.al. [6] have shown that the inner working of the CNN can be projected back to the image space in a comprehensible way to a human expert. We build on the work of [5] and present an approach to understand the decision making process of these networks through visualizing the information used as a part of this process. Our approach is based on fractionally strided convolutional technique [5], which we apply to the anatomy classification problem using X-ray images [2]. However, rather than examining the model over the whole dataset and trying to understand the sensitivity of the model to the data, we examine the model's response to individual data points. We also found that existing methods that present different saliency maps of the sensitivity of the model's output still not provide a clear representation that can be used to communicate with the experts. We based our approach on visualizing the maximally activated feature maps from the last convolutional layer in an overlayed image. This depiction provided the most informative and effective way to communicate the information from the image used in the decision process of the model. Furthermore, we compared this information to medically relevant landmarks in the images such as anatomical features that an

*Part of this work was also done while affiliated to University of Waterloo, ON, Canada - N2L 3G1
†Part of this work was also done while affiliated to Philips Research, Eindhoven Netherlands

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

expert would use to identify an organ. In this comparison, we found that shallow models that do not have sufficient capacity fail to use relevant landmarks. Additionally, we found that even deep models that generally perform well on test data do not necessarily use accurate landmarks. Finally, we show that adjusting training and augmentation hyperparameters based on the insight from visualization leads to models that use medically relevant landmarks while attaining superior performance on test data and give indication of robustness in terms of generalization.

2 Approach

In order to understand the decision making process of deep CNNs and to construct an informed approach to designing models, we build three different deep CNN models with different architectures and hyper-parameters: shallow CNN, deeper CNN without data augmentation and deeper CNN with data augmentation inspired by the work of Razavian et al. [3]. The network architecture for each model is described in Appendix (Section 5). After successfully training the above mentioned multiple networks, we examine which part of a particular input image from an anatomy class, particularly the spatially distributed information, is used in the decision processing of CNN. It is done by visualizing the top n most *activated* activation filters of the last convolutional layer in the above described models. Top n filters are used to visualize the parts of the input image that the network considers *important*. The visualization itself is done by projecting the top filter activations back to image space. The back projection to input space is achieved by using the fractionally strided convolutional technique [5] through another parallel network. The parallel network is constructed similarly to the one being analyzed but with transposed convolution filters and switches for un-pooling. This leads to unsupervised construction of hierarchical image presentation in the visual domain for any filter in a particular layer of the network.

Next, we then examine the correlation of those regions obtained through visualization with identified regions and shapes of image landmarks that are mentioned in the medical radiology literature. With the qualitative assessment, we can establish that the same landmarks that are described in the medical image literature are also used by the deep CNN. For example, we observe that the particular outlines of bones are used to detect the organ in the image rather than some background information. We use this to guide the decisions for the model architecture and training. We can furthermore use this method to detect biases in the models. In certain examples of mis-classification (Fig. 4) , we can observe that the information used for making decisions is part of an artifact rather than the object in the image. This understanding can inform us about the possible adjustments to the pre-processing of data augmentation procedures needed to remove the bias from the model.

3 Experiments and Results

To visualize and understand the decision working of a deep neural network, we used anatomy classification from X-ray images as an example use-case. To train our three different convolutional neural networks, radiographs from the ImageClef 2009 - Medical Image Annotation task [3] was used. This data set consists of a wide range of X-rays images for clinical routine, along with a detailed anatomical classification. For uniform training without any bias, we removed the hierarchical class representation and removed the classes consisting of less than 50 examples. Using this, we ended up with 24 unique classes e.g. foot, knee, hand, cranium,thoracic spine etc from the full body anatomy.

For training the three network described in Section 2, we resized the images to 224×224. For evaluation, we divided the ImageClef dataset (14676) images into randomly selected training and test set with 90 % and 10 % data respectively. For the third (deeper) network specifically, we used various data augmentation techniques ranging from cropping, rotation, translation, shearing, stretching and flipping. We train the three networks for all the 24 classes simultaneously. The results obtained by training the three models are shown in Table 1.

We visualized the internal activations of the models on test data. More particularly, we combined the visualization of the top $n = 25$ filter responses from the last convolutional layer and overlay them on the original image. In this way we construct the focused heats maps that can be easily examined by a human expert. The $n = 25$ was chosen empirically as it produced heat maps closer to the anatomical

[3] http://www.imageclef.org/2009/medanno

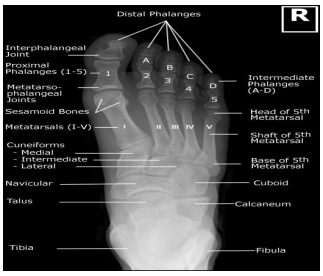

(a) Anatomy description of foot found in literature, unique bone structures pertaining to the class are indicated.

(b) Foot X-ray image from ImageClef dataset.

(c) Heat map from top 25 filters from last conv layer of network overlayed on original image.

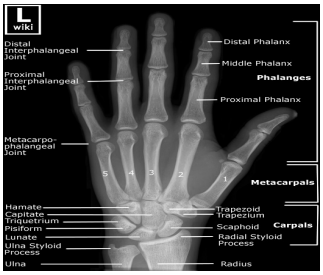

(d) Anatomy description of hand found in literature, unique bone structures pertaining to the class are indicated.

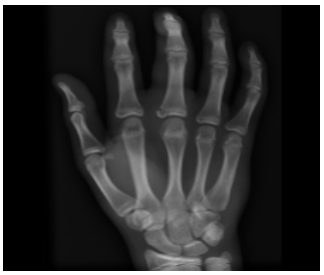

(e) Hand X-ray image from ImageClef dataset.

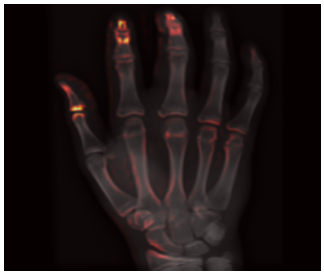

(f) Heat map from top 25 filters from last conv layer of network overlayed on original image.

Figure 1: Figure showing the correspondence between anatomical description of found in literature that are used by human experts ((a) & (d)) and the heat maps overlayed on the original images ((b) & (e)) from the last conv layer of the deeper network with data augmentation ((c) & (f)) for the foot and hand class. It can be observed that the deeper neural network uses the same landmarks as a human experts for anatomy classification. Best viewed in color.

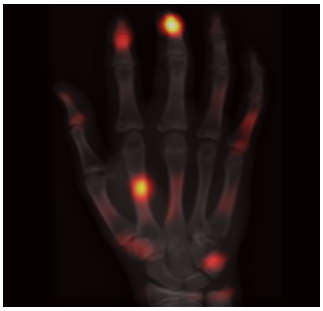

(a) Heat map from top 25 filters from last conv layer of network for the hand class.

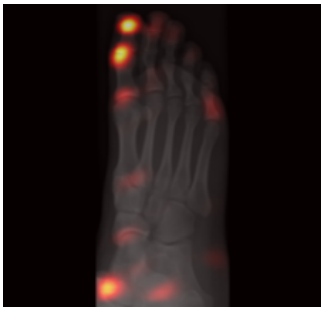

(b) Heat map from top 25 filters from last conv layer of network for the foot class.

Figure 2: Figure showing heat maps (a) & (b) overlayed on the original images from the last conv layer of the deeper network with no data augmentation overlayed on the original images for the foot and hand class. It can be observed that this network fails to use the same landmarks as a human experts for anatomy classification, as shown in Fig. 1 (a) & (c). Best viewed in color.

Table 1: Results: Accuracy in percent for three different networks trained on the imageClef 2009 annotation task

Shallow Network	**Deeper Network**	**Deeper Network+data aug**
71.1	90.36	**95.62**

Figure 3: Figure showing the focus area of the top 5 last conv layer of the shallow network. For clarity, instead of top 25 only top 5 filters are shown separately. It is evident that the network doesn't learn any medically relevant landmark. Best viewed in color.

landmarks with least number of filters. The results thus obtained are shown in Fig. 1, 2 and 3 for foot and hand classes from ImageClef dataset.

In Fig. 1 we show a correspondence between the obtained heat maps and the anatomical landmarks from medical literature [4]. More particularly for the foot image we can observe that the edges of the metatarsals' shaft has been used together with the distal phalangies, navicular, cuboid, tibia, and fibula. Similarly for the hand, three of the distal phalanxes, many of the heads of joints, metacarpals' shafts as well certain carpals. In contrast to this in Fig 3 and Fig 2 we can see that the shallow and deep network trained without specific data augmentation fails to learn such specific landmarks. These models use broader ranges that are clearly not as specific as the information used in the first model. From the above visual results [5] as well as the performance of the final model we come to the conclusion that sufficiently deep neural network model can be successfully trained to use the same medical landmarks as a human expert while attaining superior performance.

4 Conclusion

We propose an approach that allows for evaluating the decision making process of CNNs. We show that the design of the model architectures for deep CNN and the training procedure does not necessarily need to be a trial-and-error process, solely focused on optimizing the test set accuracy. Through visualization we managed to incorporate domain knowledge and overall managed to achieve a much more informed decision process, which finally resulted in a model with superior performance. This approach is applicable to many different image analysis applications of deep learning that are unable to easily leverage the potentially large amount of available domain knowledge. Furthermore, visually understanding the information involved in the model decision allows for more confidence in its performance on unseen data.

References

[1] Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In *2015 IEEE conference on computer vision and pattern recognition (CVPR)*, pages 5188–5196. IEEE, 2015.

[2] Vlado Menkovski, Zharko Aleksovski, Axel Saalbach, and Hannes Nickisch. Can pretrained neural networks detect anatomy? *arXiv preprint arXiv:1512.05986*, 2015.

[3] Ali Sharif Razavian, Josephine Sullivan, Atsuto Maki, and Stefan Carlsson. A baseline for visual instance retrieval with deep convolutional networks. *arXiv preprint arXiv:1412.6574*, 2014.

[4] Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[4] `http://www.meddean.luc.edu/lumen/meded/radio/curriculum/bones/Strcture_Bone_teach_f.htm`

[5] We obtained similar results for the other classes as well, but due to space constraint only two classes results are shown.

[5] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision*, pages 818–833. Springer, 2014.

[6] Luisa M Zintgraf, Taco S Cohen, and Max Welling. A new method to visualize deep neural networks. *arXiv preprint arXiv:1603.02518*, 2016.

5 Appendix

Table 2: Architecture of Shallow Network

Conv Layer	(3x3, 32x)
Conv Layer	(3x3, 64x)
MaxPool Layer + Dropout	(2x2, 2x2 stride)
FC Layer + Dropout	(128, 0.25p)
Softmax	(24)

Table 3: Architecture of Deeper Network without data augmentation

Conv Layer	(3x3, 32x)
Conv Layer	(3x3, 16x)
MaxPool Layer	(2x2, 2x2 stride)
Conv Layer	(3x3, 64x)
Conv Layer	(3x3, 32x)
MaxPool Layer	(2x2, 2x2 stride)
Conv Layer	(3x3, 128x)
Conv Layer	(3x3, 128x)
Conv Layer	(3x3, 64x)
MaxPool Layer	(2x2, 2x2 stride)
Conv Layer	(3x3, 256x)
Conv Layer	(3x3, 256x)
Conv Layer	(3x3, 128x)
MaxPool Layer	(2x2, 2x2 stride)
FC Layer	(128x)
FC Layer	(128x)
Softmax	(24)

Table 4: Architecture of Deeper Network with data augmentation

Conv Layer + BN	(3x3, 32x)
Conv Layer + BN	(3x3, 16x)
MaxPool Layer	(2x2, 2x2 stride)
Conv Layer + BN	(3x3, 64x)
Conv Layer + BN	(3x3, 32x)
MaxPool Layer	(2x2, 2x2 stride)
Conv Layer + BN	(3x3, 128x)
Conv Layer + BN	(3x3, 128x)
Conv Layer + BN	(3x3, 64x)
MaxPool Layer	(2x2, 2x2 stride)
Conv Layer + BN	(3x3, 256x)
Conv Layer + BN	(3x3, 256x)
Conv Layer + BN	(3x3, 128x)
MaxPool Layer+Dropout	(2x2, 2x2 stride, 0.5p)
FC Layer + Dropout	(128x, 0.5p)
FC Layer + Dropout	(128x, 0.5p)
Softmax	(24)

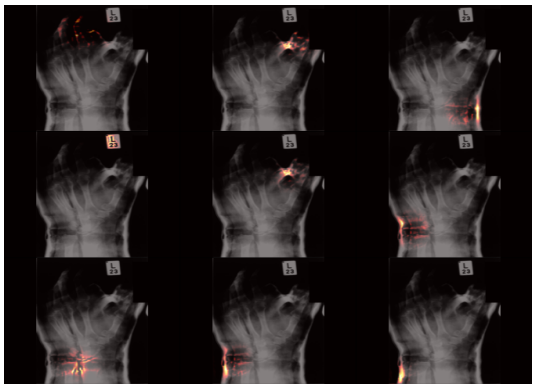

Figure 4: Figure showing the results from the last conv layer of the deeper network with augmentation for an example from the hand class mis-classified as cranium. From the figure, it is evident that the top 9 most activated filters are focusing on wrong information present in the signal. Best viewed in color.